



A Unified Algebraic and Logic-Based Framework **towards Safe Routing Implementations**

Boon Loo
TRUSTEES OF THE UNIVERSITY OF PENNSYLVANIA

08/13/2015
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

| | | | | | |
|--|--------------------|--------------------------------|-----------------------------------|--|--|
| REPORT DOCUMENTATION PAGE | | | | Form Approved OMB No. 0704-0188 | |
| <p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.</p> | | | | | |
| 1. REPORT DATE (DD-MM-YYYY) 15-08-2015 | | 2. REPORT TYPE FINAL | | 3. DATES COVERED (From - To) 6/15/2012-6/15/2015 | |
| 4. TITLE AND SUBTITLE "(YIP) A Unified Algebraic & Logic - Based Framework Towards Safe Routing Implementations" | | | | 5a. CONTRACT NUMBER | |
| | | | | 5b. GRANT NUMBER FA9550-12-1-0327 | |
| | | | | 5c. PROGRAM ELEMENT NUMBER | |
| 6. AUTHOR(S) BOON Thau LOO, Ph.D., Principal Investigator | | | | 5d. PROJECT NUMBER | |
| | | | | 5e. TASK NUMBER | |
| | | | | 5f. WORK UNIT NUMBER | |
| 7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Trustees of the University of Pennsylvania Office of Research Services 3451 Walnut Street Philadelphia PA 19104-6205 | | | | 8. PERFORMING ORGANIZATION REPORT NUMBER FINAL | |
| 9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) USAF, AFRL DUNS 143574726 AF OFFICE OF SCIENTIFIC RESEARCH 875 N. RANDOLPH ST. ROOM 3112 ARLINGTON VA 22203 | | | | 10. SPONSOR/MONITOR'S ACRONYM(S) FA9550 | |
| | | | | 11. SPONSOR/MONITOR'S REPORT NUMBER(S) | |
| 12. DISTRIBUTION/AVAILABILITY STATEMENT Program Manager; technicalreports@afosr.mil; unlimited DISTRIBUTION A | | | | | |
| 13. SUPPLEMENTARY NOTES | | | | | |
| 14. ABSTRACT This is the final report for the AFOSR Young Investigator award. The report summarizes work done in the final third year in the areas of software-defined networking use cases. We next provide research highlights spanning all three years, and discuss manpower training and a list of publications. | | | | | |
| 15. SUBJECT TERMS Formal methods, networking, logic, routing algebra, software defined networking | | | | | |
| 16. SECURITY CLASSIFICATION OF: | | | 17. LIMITATION OF ABSTRACT | 18. NUMBER OF PAGES | 19a. NAME OF RESPONSIBLE PERSON |
| a. REPORT | b. ABSTRACT | c. THIS PAGE | | | 19b. TELEPHONE NUMBER (Include area code) |

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATES COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5d. PROJECT NUMBER. Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

Contract/Grant Title: A Unified Algebraic & Logic-Based Framework Towards Safe Routing Implementations

Contract/Grant #: FA9550-12-1-0327

Principle Investigator: Boon Thau Loo

Reporting Period: 15 Jun 2014 to 14 Jun 2015 (Final Report)

1.0 Highlights in the previous year (15 Jun 2015 – 14 Jun 2015)

In the past year, we have focused on two aspects of work. First, we applied concepts developed in the first two years in the domain of Software-defined Networks (SDN). We developed a declarative platform for implementing SDN protocols using declarative networking programs that can be automatically verified for correctness. Second, we have also developed a declarative platform for automatically synthesizing SDN protocols from example scenarios.

Declarative network verification. Networks are complex systems that unfortunately are ridden with errors. Such errors can lead to disruption of services, which may have grave consequences. Verification of networks is key to eliminating errors and building robust networks. Over the past year, we propose an approach to verify networks using declarative networking, where networks are specified in Network Datalog, a declarative language. We focus on analyzing safety properties. We have developed a technique to statically analyze NDlog programs: first, we build a dependency graph of the predicates of NDlog programs; then, we build a summary data structure called a derivation pool to represent all possible derivations and their associated constraints for predicates in the program; finally, properties specified in first-order logic are checked on the data structure with the help of the SMT solver Z3. We build a prototype tool and demonstrate the effectiveness of the tool in validating and debugging several SDN applications.

Example-based SDN synthesis. Recent emergence of software-defined networks offers an opportunity to design domain-specific programming abstractions aimed at network operators. We propose scenario-based programming, a framework that allows network operators to program network policies by describing representative example behaviors. Given these scenarios, our synthesis algorithm automatically infers the controller state that needs to be maintained along with the flowtable rules to process network events and update state. We have developed the NetEgg scenario-programming tool, which can execute the generated policy implementation on top of a centralized controller, but also automatically infers rules that can be pushed to switches to improve throughput. We study a range of policies considered in the literature and report our experience regarding specifying these policies using scenarios. We evaluate NetEgg based on the computational requirements of our synthesis algorithm as well as the overhead introduced by the generated policy implementation. Our results show that our synthesis algorithm can generate policy implementations in seconds, and the automatically generated policy implementations have performance comparable to their handcrafted implementations.

2.0 Summary of research (15 Jun 2012 – 14 Jun 2015)

Over the past three years, we carried out several areas of work that aims to unify algebraic and logic-based framework for Safe Routing Implementations. The original basis of our work started from the Formally Safe Routing (FSR) toolkit. The FSR toolkit attempts to bridge the gap in formal reasoning and distributed implementations, by unifying research in routing algebras with recent advances in declarative networking to produce provably correct distributed implementations. Specifically, FSR automates the process of analyzing routing configurations expressed in algebra for safety (i.e. convergence) using SMT solvers, and automatically compiles routing algebra into declarative routing implementations.

Based on the FSR toolkit, we have worked on the following areas of work:

BGPVerif – To scale up formal analysis on Internet routing systems, we propose a novel scalability technique, called network reduction, by exploring the networking problem space. Based on network reduction, we build a formal analysis toolkit BGPVerif, which enables networking researchers to study and analyze large BGP (Border Gateway Protocol) systems in a sound and automatic fashion. Central to BGPVerif is a network reduction engine for rewriting policy configurations into smaller and simpler forms, and an automatic analyzer for performing analysis on the BGP instance. In network reduction, we provide two types of reduction rules that transform policy configurations by merging duplicate and complementary router configurations to simplify analysis. We show that the network reductions are sound, dual of each other and are locally complete. The reductions are also computationally attractive, requiring only local configuration information and modification. We have developed a prototype of network reduction and demonstrated that it is applicable on various BGP systems and enables significant savings in analysis time. In addition to making possible safety analysis on large networks that would otherwise not complete within reasonable time, network reduction is also a useful tool for discovering possible redundancies in BGP systems. BGPVerif appeared at Infocom 2014, and together with the earlier FSR work, formed the basis of Anduo Wang's PhD dissertation work.

SDN verification and synthesis -- We explored a formal synthesis approach for managing SDNs. With the tremendous growth of the Internet and the emerging SDN infrastructure, there is an increasing need for rigorous and scalable network management methods and tool support. We explore a synthesis approach for managing software-defined networks. We formulate the construction of network control logic as a reactive synthesis problem, which is solvable with existing synthesis tools. The key idea is to synthesize a strategy that manages control logic in response to network changes while satisfying some network-wide specification. Finally, we investigate network abstractions for scalability. For large networks, instead of synthesizing control logic directly, we use its abstraction-a smaller network that simulates its behavior-for synthesis, and then implement the synthesized control on the original network while preserving the

correctness. By using the so-called simulation relations, we also prove the soundness of this abstraction-based synthesis approach.

An initial version of this work appeared in the WRiPE 2013. We have since extended this to the scenario-based programming platform called NetEgg described earlier in Section 1 in this report. NetEgg appeared in ACM HotNets 2014, and a paper is under submission to CoNEXT 2015.

FixRoute toolkit – The FixRoute toolkit extends FSR by incorporating numerical optimizations, to ensure safe routing implementations that meet traffic engineering goals. FixRoute is motivated by the fact that the performance of networks that use the Internet Protocol is sensitive to precise configuration of many low-level parameters on each network device. These settings govern the action of dynamic routing protocols, which direct the flow of traffic; in order to ensure that these dynamic protocols all converge to produce some 'optimal' flow, each parameter must be set correctly. Multiple conflicting optimization objectives, non-determinism, and the need to reason about different failure scenarios make the task particularly complicated. We present a fast and flexible approach for the analysis of a number of such management tasks presented in the context of BGP routing. The idea is to combine logical satisfiability criteria with traditional numerical optimization, to reach a desired traffic flow outcome subject to given constraints on the routing process. The method can then be used to probe parameter sensitivity, trade-offs in the selection of optimization goals, resilience to failure, and so forth. The theory is underpinned by a rigorous abstraction of the convergence of distributed asynchronous message-passing protocols, and is therefore generalizable to other scenarios. Our resulting hybrid engine is faster than either purely logical or purely numeric alternatives, making it potentially feasible for interactive production use. An earlier version of FixRoute (called Route Sherperd) was demonstrated at SIGCOMM 2012 and a paper on FixRoute is currently under submission at Infocom 2016.

Declarative network verification – Our fourth area of focus is on generating provably correct network implementations using declarative networking techniques. See Section 1.0 for details. Work on this area appeared in PPDP 2015.

Declarative secure network verification – Our fifth area of work is in extending the declarative network verification work in the security arena. The Internet, as it stands today, is highly vulnerable to attacks. However, little has been done to understand and verify the formal security guarantees of proposed secure inter-domain routing protocols, such as Secure BGP (S-BGP). We develop a sound program logic for SANDLog -- a declarative specification language for secure routing protocols -- for verifying properties of these protocols. We prove invariant properties of SANDLog programs that run in an adversarial environment. As a step towards automated verification, we implement a verification condition generator (VCGen) to automatically extract proof obligations. VCGen is integrated into a compiler for SANDLog that can generate executable protocol

implementations; and thus, both verification and empirical evaluation of secure routing protocols can be carried out in this unified framework. To validate our framework, we encoded several proposed secure routing mechanisms in SANDLog, verified variants of path authenticity properties by manually discharging the generated verification conditions in Coq, and generated executable code based on SANDLog specification and ran the code in simulation. Work on this area appeared in FORTE 2014.

3.0 Manpower training

We have graduated several Ph.D. students and postdocs that work on topics related to this grant. They include

- Anduo Wang (Ph.D. 2013, UIUC postdoc, joining Temple University as Assistant Professor in January 2015)
- Dong Lin (Ph.D. 2015, joined Google)

In addition, current 4th year Ph.D. student Chen Chen also worked on topics related to this grant, and will graduate by summer 2016.

4.0 Main publications generated (15 Jun 2012 – 14 Jun 2015)

Automated Verification of Safety Properties in Declarative Networking Programs. Chen Chen, Lay Kuan Loh, Limin Jia, Wenchao Zhou, and Boon Thau Loo. 17th International ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming (PPDP), July, 2015.

A Scalable Multi-Datacenter Layer-2 Network Architecture. Chen Chen, Changbin Liu, Pingkai Liu, Boon Thau Loo, and Ling Ding. Symposium on SDN Research (SOSR), 2015.

Private and Verifiable Interdomain Routing Decisions. Mingchen Zhao, Wenchao Zhou, Alexander J. T. Gurney, Andreas Haeberlen, Micah Sherr, and Boon Thau Loo. IEEE/ACM Transactions on Networking (ToN), 2015.

NetEgg: Programming Network Policies by Examples. Yifei Yuan, Rajeev Alur, and Boon Thau Loo. 13th ACM Workshop on Hot Topics in Networks (HotNets-XIV), 2014.

Diagnosing Missing Events in Distributed Systems with Negative Provenance. Yang Wu, Mingchen Zhao, Andreas Haeberlen, Wenchao Zhou, and Boon Thau Loo. ACM SIGCOMM Conference on Data Communication, 2014.

Program Logic for Verifying Secure Routing Protocols. Chen Chen, Limin Jia, Hao Xu, Cheng Luo, Wenchao Zhou and Boon Thau Loo. A 34th IFIP International Conference on Formal Techniques for Distributed Objects, Components and Systems (FORTE). Berlin, Germany, 2014.

A Reduction-based Approach Towards Scaling Up Formal Analysis of Internet Configurations. Anduo Wang, Alexander Gurney, Xianglong Han, Jinyan Cao, Boon Thau Loo, Carolyn Talcott, and Andre Scedrov. 33rd Annual IEEE International Conference on Computer Communications (INFOCOM). Toronto, Canada, 2014.

Automated Synthesis of Reactive Controllers for Software-Defined Networks. Anduo Wang, Salar Moarref, Ufuk Topcu, Boon Thau Loo and Andre Scedrov. 3rd International Workshop on Rigorous Protocol Engineering (WRiPE), 2013.

Declarative Platform for High-Performance Network Traffic Analytics. Harjot Gill, Dong Lin, Cam Nguyen, Tanveer Gill, and Boon Thau Loo. Cluster Computing journal (Special Issue on selected best papers of HPDC 2013), 2014.

Impact of Path Selection and Scheduling Policies on MPTCP Performance. Behnaz Arzani, Alexander Gurney, Shuotian Cheng, Roch Guerin and Boon Thau Loo. 4th International Workshop on Protocols and Applications with Multi-Homing Support (PAMS), 2014.

A Formal Framework for Secure Routing Protocols. Chen Chen, Limin Jia, Hao Xu, Cheng Luo, Wenchao Zhou, and Boon Thau Loo. USENIX Symposium on Networked Systems Design and Implementation (NSDI) (demonstration), 2013.

Reduction-based Analysis of BGP Systems with BGPVerif. Anduo Wang, Alexander J.T. Gurney, Xianglong Han, Jinyan Cao, Carolyn Talcott, Boon Thau Loo, and Andre Scedrov. ACM SIGCOMM Conference on Data Communication (demonstration), Helsinki, Finland, Aug, 2012.

Route Shepherd: Stability Hints for the Control Plane. Alexander J.T. Gurney, Xianglong Han, Yang Li, and Boon Thau Loo. ACM SIGCOMM Conference on Data Communication (demonstration), Helsinki, Finland, Aug, 2012.

A Formal Framework for Secure Routing Protocols. Chen Chen, Limin Jia, Hao Xu, Cheng Luo, Wenchao Zhou and Boon Thau Loo. Workshop on Foundations of Computer Security (FCS), co-located with CSF, New Orleans, Louisiana, June 2013.

Reduction-based Security Analysis of Internet Routing Protocols. Chen Chen, Limin Jia, Boon Thau Loo, and Wenchao Zhou. 2nd International Workshop on Rigorous Protocol Engineering (WRiPE), co-located with ICNP 2012, Oct 2012.

Brief Announcement: A Calculus of Policy-Based Routing Systems. Anduo Wang, Carolyn Talcott, Alexander J.T. Gurney, Boon Thau Loo, and Andre Scedrov. 31st Annual ACM SIGACT-SIGOPS Symposium on Principles of Distributed Computing (PODC), July, 2012.

1.

1. Report Type

Final Report

Primary Contact E-mail**Contact email if there is a problem with the report.**

boonloo@cis.upenn.edu

Primary Contact Phone Number**Contact phone number if there is a problem with the report**

215-898-3323

Organization / Institution name

University of Pennsylvania

Grant/Contract Title**The full title of the funded effort.**

(YIP) A Unified Algebraic & Logic - Based Framework Towards Safe Routing Implementations

Grant/Contract Number**AFOSR assigned control number. It must begin with "FA9550" or "F49620" or "FA2386".**

FA9550-12-1-0327

Principal Investigator Name**The full name of the principal investigator on the grant or contract.**

Boon Thau Loo

Program Manager**The AFOSR Program Manager currently assigned to the award**

James Lawton and Tristan Nguyen

Reporting Period Start Date

06/15/2012

Reporting Period End Date

06/14/2015

Abstract

In the past year, we have focused on two aspects of work. First, we applied concepts developed in the first two years in the domain of Software-defined Networks (SDN). We developed a declarative platform for implementing SDN protocols using declarative networking programs that can be automatically verified for correctness. We build a prototype tool and demonstrate the effectiveness of the tool in validating and debugging several SDN applications.

Second, we have also developed a declarative platform for automatically synthesizing SDN protocols from example scenarios. Our results show that our synthesis algorithm can generate policy implementations in seconds, and the automatically generated policy implementations have performance comparable to their handcrafted implementations.

Distribution Statement**This is block 12 on the SF298 form.**

Distribution A - Approved for Public Release

Explanation for Distribution Statement**If this is not approved for public release, please provide a short explanation. E.g., contains proprietary information.**

DISTRIBUTION A: Distribution approved for public release

SF298 Form

Please attach your [SF298](#) form. A blank SF298 can be found [here](#). Please do not password protect or secure the PDF. The maximum file size for an SF298 is 50MB.

[PIY BTL AFD-070820-035.pdf](#)

Upload the Report Document. File must be a PDF. Please do not password protect or secure the PDF . The maximum file size for the Report Document is 50MB.

[final-report.pdf](#)

Upload a Report Document, if any. The maximum file size for the Report Document is 50MB.

Archival Publications (published) during reporting period:

See report.

Changes in research objectives (if any):

None

Change in AFOSR Program Manager, if any:

None

Extensions granted or milestones slipped, if any:

None.

AFOSR LRIR Number

LRIR Title

Reporting Period

Laboratory Task Manager

Program Officer

Research Objectives

Technical Summary

Funding Summary by Cost Category (by FY, \$K)

| | Starting FY | FY+1 | FY+2 |
|----------------------|-------------|------|------|
| Salary | | | |
| Equipment/Facilities | | | |
| Supplies | | | |
| Total | | | |

Report Document

Report Document - Text Analysis

Report Document - Text Analysis

Appendix Documents

2. Thank You

E-mail user

Aug 12, 2015 14:28:13 Success: Email Sent to: boonloo@cis.upenn.edu